

WHAT IS CLAIMED IS:

1. At least one computer-readable medium having computer program instructions embodied therein for interacting with at least one knowledge base, the at least one knowledge base having data stored therein representing first knowledge about a plurality of objects using a plurality of relationships between pairs of the objects, the computer program instructions comprising:

first computer program instructions operable to interact with the at least one knowledge base and infer second knowledge not represented in the at least one knowledge base from the first knowledge; and

second computer program instructions operable to respond to a query having a predetermined format using at least one of the first knowledge and the second knowledge.

2. The at least one computer-readable medium of claim 1 wherein the first computer program instructions are operable to infer the second knowledge with reference to the relationships.

3. The at least one computer-readable medium of claim 1 wherein at least a portion of the first and second computer program instructions are part of a plug-in application for use with a client browser application.

4. The at least one computer-readable medium of claim 1 wherein at least a portion of the first and second computer program instructions are part of a server-side application.

5. The at least one computer-readable medium of claim 1 wherein the computer program instructions further comprise third computer program instructions operable to generate the query.

5 6. The at least one computer-readable medium of claim 5 wherein the third computer program instructions are operable to generate the query in response to query data representing a natural language question.

10 7. The at least one computer-readable medium of claim 6 wherein the natural language question is entered as a text string into a URL window in a browser environment, the third computer program instructions further being operable to determine whether the text string is a URL or the query data.

15 8. The at least one computer-readable medium of claim 5 wherein the third computer program instructions are operable to translate a natural language question into the query by generating at least one sub-string from a string of text corresponding to the natural language question, and selecting at least one of a plurality of query template components corresponding to the at least one sub-string to generate the query.

20 9. The at least one computer-readable medium of claim 8 wherein the third computer program instructions are further operable to substitute the selected query template components into variables associated with a query generator execution of which generates the query.

10. The at least one computer-readable medium of claim 8 further comprising fourth computer program instructions operable to generate a query template including the plurality of query template components by designating a plurality of predefined text sub-strings and a plurality of variables to which the predefined text sub-strings correspond, and
5 by defining a query generator with respect to the variables, the query generator being operable to generate the query from selected ones of the predefined text sub-strings being substituted for the corresponding variables.

11. The at least one computer-readable medium of claim 5 wherein the third
10 computer program instructions are further operable to generate a plurality of query candidates including the query.

12. The at least one computer-readable medium of claim 11 wherein the third
15 computer program instructions are further operable to reject selected ones of the plurality of query candidates with reference to the knowledge base.

13. The at least one computer-readable medium of claim 11 wherein the third
computer program instructions are further operable to present the plurality of query candidates for selection of the query.

14. The at least one computer-readable medium of claim 1 wherein the computer
program instructions further comprise third computer program instructions operable to
generate a natural language response to the query with reference to selected ones of the
plurality of objects.

15. The at least one computer-readable medium of claim 14 wherein the third computer program instructions are operable to generate the natural language response with reference to translation objects associated with the selected objects in the knowledge base.

5 16. The at least one computer-readable medium of claim 15 wherein the translation objects comprise common translation objects and unique translation objects, the common translation objects relating an associated one of the selected objects with a common translation for the associated selected object, and the unique translation objects relating an associated one of the selected objects with a unique translation for the associated object.

10 17. The at least one computer-readable medium of claim 1 wherein the first computer program instructions are operable to implement a plurality of inference generators each of which is operable to generate facts not represented in the knowledge base with reference to at least one of the plurality of objects and at least one potential fact associated with the inference generator.

15 18. The at least one computer-readable medium of claim 17 wherein some of the plurality of inference generators are operable to generate facts with reference to the associated objects, the relationships corresponding to the associated objects, and the associated potential facts, and without reference to any processes external to the inference generator.

20 19. The at least one computer-readable medium of claim 17 wherein some of the plurality of inference generators are operable to generate facts with reference to the associated objects, the relationships corresponding to the associated objects, and the

associated potential facts, and only after performing at least one additional operation using a at least one process external to the inference generator.

20. The at least one computer-readable medium of claim 1 wherein the data
5 representing the first knowledge comprises a plurality of fact representations each corresponding to a relationship between a pair of the plurality of objects, each fact representation including a first object name corresponding to a first one of the plurality of objects, a second object name corresponding to a second one of the plurality of objects, and a relation object name identifying a relation object which defines the relationship between
10 the first and second objects, and wherein the first computer program instructions are operable to infer the second knowledge with reference to the relationships.

21. A computer-implemented method for interacting with at least one knowledge
base, the at least one knowledge base having data stored therein representing first knowledge
15 about a plurality of objects using a plurality of relationships between pairs of the objects, the method comprising:

interacting with the at least one knowledge base and inferring second knowledge not represented in the at least one knowledge base from the first knowledge; and

responding to a query having a predetermined format using at least one of the first
20 knowledge and the second knowledge.

22. The method of claim 21 wherein the second knowledge is inferred with reference to the relationships.

23. The method of claim 21 further comprising generating the query.

24. The method of claim 23 wherein the query is generated in response to query data representing a natural language question.

5 25. The method of claim 24 wherein the natural language question is entered as a text string into a URL window in a browser environment, the method further comprising determining whether the text string is a URL or the query data.

0990188.112101
26. The method of claim 23 further comprising translating a natural language
10 question into the query by generating at least one sub-string from a string of text corresponding to the natural language question, and selecting at least one of a plurality of query template components corresponding to the at least one sub-string to generate the query.

15 27. The method of claim 26 further comprising substituting the selected query template components into variables associated with a query generator execution of which generates the query.

20 28. The method of claim 26 further comprising generating a query template including the plurality of query template components by designating a plurality of predefined text sub-strings and a plurality of variables to which the predefined text sub-strings correspond, and by defining a query generator with respect to the variables, the query generator being operable to generate the query from selected ones of the predefined text sub-strings being substituted for the corresponding variables.

5

10

15

20

35. The method of claim 21 further comprising implementing a plurality of inference generators each of which is operable to generate facts not represented in the

knowledge base with reference to at least one of the plurality of objects and at least one potential fact associated with the inference generator.

36. The method of claim 35 wherein some of the plurality of inference generators
5 are operable to generate facts with reference to the associated objects, the relationships corresponding to the associated objects, and the associated potential facts, and without reference to any processes external to the inference generator.

37. The method of claim 35 wherein some of the plurality of inference generators
10 are operable to generate facts with reference to the associated objects, the relationships corresponding to the associated objects, and the associated potential facts, and only after performing at least one additional operation using a at least one process external to the inference generator.

38. The method of claim 21 wherein the data representing the first knowledge
15 comprises a plurality of fact representations each corresponding to a relationship between a pair of the plurality of objects, each fact representation including a first object name corresponding to a first one of the plurality of objects, a second object name corresponding to a second one of the plurality of objects, and a relation object name identifying a relation
20 object which defines the relationship between the first and second objects, and wherein the method further comprises inferring the second knowledge with reference to the relationships.

39. A system for storing and retrieving information from a knowledge base,
comprising:

at least one memory having data corresponding to the knowledge base stored therein representing first knowledge about a plurality of objects using a plurality of relationships between pairs of the objects; and

at least one central processing unit operable in association with the at least one memory to interact with the at least one knowledge base and infer second knowledge not represented in the at least one knowledge base from the first knowledge, and respond to a query having a predetermined format using at least one of the first knowledge and the second knowledge.

40. The system of claim 39 wherein the at least one memory and the at least one central processing unit are connected via a memory bus.

41. The system of claim 39 wherein the at least one memory and the at least one central processing unit are connected via a local area network.

42. The system of claim 39 wherein the at least one memory and the at least one central processing unit are connected via a wide area network.

43. At least one computer readable medium having data stored therein corresponding to first knowledge about a plurality of objects comprising a plurality of fact representations each corresponding to a relationship between a pair of the plurality of objects, each fact representation including a fact name identifying the fact representation, a first object name corresponding to a first one of the plurality of objects, a second object name corresponding to a second one of the plurality of objects, and a relation object name

identifying a relation object which defines the relationship between the first and second objects.

44. The at least one computer-readable medium of claim 43 wherein each of the plurality of objects has a unique object name associated therewith.

45. The at least one computer-readable medium of claim 44 wherein each unique object name comprises a natural language noun phrase.

46. The at least one computer-readable medium of claim 43 wherein the plurality of objects includes a plurality of parametered objects, each parametered object representing a class of objects having at least one commonality.

47. The at least one computer-readable medium of claim 46 wherein each parametered object is represented by a class name corresponding to a class, and at least one parameter by which specific instances of the class are defined.

48. The at least one computer-readable medium of claim 43 wherein the plurality of fact representations includes temporal fact representations each of which refers to at least one other of the fact representations and represents at least one temporal restraint relating to the at least one other of the fact representations.

49. The at least one computer-readable medium of claim 43 wherein the plurality of fact representations represents only facts which are currently true.

0950100-1-12101

50. The at least one computer-readable medium of claim 43 wherein selected ones of the plurality of fact representations assert a negative relationship between the first and second objects.

5 51. The at least one computer-readable medium of claim 43 wherein each of selected ones of the plurality of objects corresponds to more than one of the fact representations.

52. A computer-implemented method comprising transmitting first data in a
10 network, the first data being derived from second data corresponding to knowledge about a plurality of objects comprising a plurality of fact representations each corresponding to a relationship between a pair of the plurality of objects, each fact representation including a first object name corresponding to a first one of the plurality of objects, a second object name corresponding to a second one of the plurality of objects, and a relation object name
15 defining the relationship between the first and second objects.

53. A computer-implemented method for adding new data to a knowledge base, the knowledge base having first data stored therein representing knowledge about a plurality of objects using a plurality of relationships between pairs of the objects, the method
20 comprising:

generating a new object name which identifies a new object uniquely in the knowledge base;

generating an output translation string which identifies the new object in a manner which facilitates understanding by a human;

identifying a principal class for the new object, the principal class being previously defined within the knowledge base;

generating at least one alternate name by which the new object may be identified; and

where the new object is not currently included within the principal class, inserting the new data in the knowledge base, the new data representing the new object name, the output translation string, the principal class, the at least one alternate name, and at least one relationship thereamong.

54. A computer system comprising

a plurality of computers linked by a network;

a knowledge base stored in at least one of the computers in which representations of factual knowledge are characterized by named relationships between named objects;

a knowledge inference module on at least one of the computers operable to infer new knowledge from the factual knowledge;

a query answering module on at least one of the computers operable to receive a query in a pre-determined format and respond to the query using at least the representations of the factual knowledge.

55. The computer system of claim 54 wherein the query answering module is operable to facilitate transmission of the factual knowledge in response to a request transmitted from a first one of the computers to a second one of the computers.

56. The computer system of claim 55 wherein the query answering module is operable to transmit more of the factual knowledge than requested thereby transmitting commercially valuable information in the network.

57. The computer system of claim 55 wherein the knowledge base includes server knowledge identifying which of the computers include specific portions of the factual knowledge thereby facilitating selection of the second computer by the first computer.

58. The computer system of claim 57 wherein the second computer is operable to employ the server knowledge to identify at least one alternate computer in response to the request from the first computer.

59. The computer system of claim 55 wherein the network is the internet.

60. The computer system of claim 59 further comprising at least one web browser operable to display the factual knowledge and the new knowledge.

61. The computer system of claim 54 further comprising a natural language translation module operable to translate a natural language question entered by a human operator into the predetermined format of the query.

62. The computer system of claim 60 wherein the web browser is operable to display a summary information screen providing general information on a named object using a portion of the factual knowledge retrieved from the knowledge base.

63. The computer system of claim 62 wherein the summary information screen includes a link corresponding to each named object which generate further summary information screens when the link is activated.

64. The computer system of claim 61 wherein the natural language translation module is operable to translate the natural language question by comparing the natural language question to pre-determined templates comprising sequences of known and unknown strings, and is further operable to use template queries to generate the query by substituting portions of the factual knowledge provided by the query answering module into at least one of the template queries.

65. The computer system of claim 64 wherein the natural language translation module is further operable to reject the query as invalid with reference to relationships between objects associated with the query where no classes of which the objects are members match any permitted classes associated with the relationships.

66. The computer system of claim 59 wherein online documents correspond to unique names and wherein the factual knowledge associates URLs for downloading the online documents with the unique names in the knowledge base.

67. The computer system of claim 54 wherein the factual knowledge includes temporal knowledge regarding when the relationships between named objects are true.

68. The computer system of claim 67 wherein the query answering module is operable to respond to the query when the query corresponds to a question requiring a human being to answer either of yes and no, a no answer indicating that a reverse relationship co-existed for an overlapping period of time with a specific relationship represented in the query.

69. The computer system of claim 54 wherein the query requests a list of a subset of the object, and the factual knowledge includes count knowledge representing a number of the objects that have a given relationship to a second object, and wherein the query
5 answering module is operable to use the count knowledge to indicate whether a response to the query is complete.

70. The computer system of claim 68 wherein the query requests a list of a subset of the object, and the factual knowledge includes count knowledge representing a number of
10 the objects that have a given relationship to a second object, and wherein the query answering module is operable to use the count knowledge to indicate whether a response to the query is complete.

71. The computer system of claim 54 wherein the named objects correspond to
15 object names which include a local identifier and a local group, and wherein each object name can only be assumed to be distinct from other object names in a same group where the local identifier is different, and wherein the system further comprises a knowledge base management module which is operable to incorporate additional knowledge into the knowledge base without first having to identify every additional object.

20 72. A computer system for storing and processing factual knowledge, comprising:

a memory

a knowledge base stored in the memory comprising the factual knowledge

25 represented with reference to groups, each group having associated therewith:

a group name for identifying the group;

a first object name for identifying a first object;

a second object name for identifying a second object; and

a relation object name for identifying a relationship between the first and second

objects;

a knowledge inference module operable to infer new knowledge not represented in the knowledge base with reference to the factual knowledge.

a query answering module operable to retrieve knowledge from the knowledge base in response to a request.

73. A computer-implemented method for storing and representing factual knowledge in a computer system having a memory, the method comprising:

representing the factual knowledge as relationships between pairs of objects, each type of relationship and each of said objects being represented by unique object names; and

generating a unique fact name for each relationship between each pair of objects, each unique fact name representing a fact representation comprising at least three of the unique object names, at least one of at least three unique object names asserting a first one of the relationships between others of the at least three unique objects names, wherein each unique fact name may be used in other fact representations.

74. A computer-implemented method for storing and representing factual knowledge in a computer system having a memory, the method comprising:

representing the factual knowledge as relationships between pairs of objects, each type of relationship and each of said objects being represented by unique object names; and

generating a unique fact name for each relationship between each pair of objects, each unique fact name representing a fact representation comprising a plurality of the unique object names at least one of which corresponds to a first one of the relationships, selected ones of the fact representations further comprising a Boolean value indicating whether the first relationship or the converse of the first relationship exists between others of the plurality of unique object names identified in the fact representation.

75. A computer-implemented method for storing and representing factual knowledge in a computer system having a memory, the method comprising:

representing the factual knowledge as relationships between pairs of objects, each type of relationship and each of said objects being represented by unique object names; and generating a unique fact name for each relationship between each pair of objects, each unique fact name representing a fact representation comprising a plurality of the unique object names at least one of which corresponds to a first one of the relationships, selected ones of the fact representations further comprising a time-period object representing a period of time that the first relationship is true.

76. A computer-implemented method for adding new knowledge to a knowledge base which includes factual knowledge, the new knowledge corresponding to a previously uncategorized object, the method comprising:

generating a new object name which uniquely identifies the previously uncategorized object within the knowledge base, and may be used in subsequently generated assertions of fact regarding the previously uncategorized object;

generating a first output translation string corresponding to a precise description of the previously uncategorized object for identifying the previously uncategorized object to a human operator;

identifying a principal class name for a class of objects of which the previously uncategorized object is a member, the principal class name previously existing in the knowledge base;

generating at least one natural language name to identify the previously uncategorized object;

where the principal class name is not associated with the new object name, associating the principal class name therewith; and

inserting into the knowledge base first data asserting a unique output relationship between the new object name and the first output translation string, and second data asserting a possible translation relationship between each of the at least one natural language names and the new object name, thereby facilitating translation of the new object name into natural language in a way which uniquely identifies the previously uncategorized object and which facilitates subsequent identification of the previously uncategorized object in the knowledge base.

77. A computer-implemented method for representing a target object in a knowledge base, the target object being a member of a large consistently-understood class of objects, the method comprising:

identifying a unique class name for the class of objects;

generating at least one unique object name which uniquely identifies the target object within the class;

combining the unique class name with the at least one unique object name using a pre-determined syntax to generate a unique target name corresponding to the target object; and

providing methods by which at least one property about the target object may be inferred with reference to the class and without reference to the unique target name.

78. A computer-implemented method of representing a query for interrogating a knowledge base comprising generating a sequence of lines using a plurality of object names and a plurality of variables, each line comprising a first object representation, a first relation object representation, and a second object representation, selected ones of the lines also comprising a fact object representation, each of the object representations comprising one of the object names and variables, wherein data retrieved from the knowledge base may be substituted into the variables in the sequence of lines such that each of the lines in the sequence corresponds to facts which may be derived from the knowledge base.

79. A computer-implemented method for responding to a query for interrogating a knowledge base, the query comprising a sequence of lines using a plurality of object names and a plurality of variables, comprising:

searching for data in the knowledge base for substitution into the variables;

where substitution of the data into the variables results in all of the lines corresponding to facts in the knowledge base, responding to the query with a yes answer;

where substitution of the data into the variables results in at least one of the lines not corresponding to any of the facts, attempting to generate a contradiction between the query and any of the facts, the contradiction comprising a reverse relationship for an overlapping period of time with respect to a first relationship in the query; and

where the contradiction is found, responding to the query with a no answer.

80. A computer-implemented method for inferring new knowledge from factual knowledge relating to a plurality of objects represented in a knowledge base, comprising:

in response to a query having at least one variable associated therewith, generating at least one new factual representation referencing the at least one variable; and

executing the query thereby causing object names corresponding to the plurality of objects to be substituted for the at least one variable in the at least one new factual representation in a manner consistent with the factual knowledge, thereby inferring the new knowledge.

81. The method of claim 80 further comprising:

determining whether the new knowledge may be inferred using the at least one
factual representation and referencing only the objects; and

where the new knowledge may not be inferred using the at least one factual representation and referencing only the objects, identifying executable program code for determining for which values of the object names the at least one fact representation is valid.

82. A translation template for translating a natural language question into a query for interrogating a knowledge base, comprising:

a plurality of substitute strings corresponding to a class including the natural language question;

a test query comprising at least one template variable; and

an output query comprising the at least one output variable;

5

10

15

20

25

84. In a system comprising a knowledge base in which knowledge is represented by relationships between selected ones of a plurality of objects and selected ones of a plurality of objects classes to which the object belong, a knowledge query module for generating responses to queries with reference to the knowledge base, the knowledge query module being operable to determine whether a first one of the objects belongs to a first one

of the object classes, and a natural language translation module operable to translate a natural language question into a plurality of possible queries for use by the knowledge query module, a computer-implemented method for rejecting invalid queries generated by the natural language translation module, the method comprising:

- 5 generating the plurality of possible queries from the natural language question using the natural language translation module; and
- for each possible query checking each corresponding relationship by:
 - identifying all object classes in the knowledge base that may be related by the corresponding relationship;
 - 10 identifying objects in the possible query that may related by the corresponding relationship;
 - interrogating said knowledge base to determine whether the objects in the possible query are members of the identified object classes; and
 - rejecting the possible query if any of the objects in the possible query are not
 - 15 members of any of the identified object classes.

85. In a system comprising a knowledge representation system and a query processing module operable to process queries and interrogate the knowledge representation system, a computer-implemented method of executing a query within a web browser,

- 20 comprising:
 - forming sequence of characters from the query which are permitted to occur within an HTML link;
 - placing a unique protocol identifier string adjacent the sequence of characters to form a pseudo-URL;

presenting the pseudo-URL as a first link within an HTML page displayed by the browser; and

in response to selection of the first link, processing the query with the query processing module and the knowledge representation system and generating a response
5 suitable for representation in the browser.

86. In a computer system comprising memory having stored therein data representative of a plurality of named objects, a plurality of assertions relating corresponding ones of the named objects to named classes of which the named objects are members, an
10 ontology relating selected ones of the named classes to more general and more specific classes, and a plurality of display templates each associated with a named class and describing generation of a general information summary about any of the named objects which is a member of the associated named class, a method for presenting information about a first one of the named objects, comprising:

15 identifying a most specific one of the named classes of which the first named object is a member, the most specific named class also having a display template associated therewith; and

generating the general information summary corresponding to the first object using the display template associated with the most specific named class.